

APPLICATION NOTE



Exploratory data structure comparisons: three new visual tools based on principal component analysis*

Anne Helby Petersen ^a, Bo Markussen^b and Karl Bang Christensen^a

^aDepartment of Public Health, University of Copenhagen, Copenhagen, Denmark; ^bDepartment of Mathematical Sciences, University of Copenhagen, Copenhagen, Denmark

ABSTRACT

Datasets are sometimes divided into distinct subsets, e.g. due to multi-center sampling, or to variations in instruments, questionnaire item ordering or mode of administration, and the data analyst then needs to assess whether a joint analysis is meaningful. The Principal Component Analysis-based Data Structure Comparisons (PCADSC) tools are three new non-parametric, visual diagnostic tools for investigating differences in structure for two subsets of a dataset through covariance matrix comparisons by use of principal component analysis. The PCADSC tools are demonstrated in a data example using European Social Survey data on psychological well-being in three countries, Denmark, Sweden, and Bulgaria. The data structures are found to be different in Denmark and Bulgaria, and thus a comparison of for example mean psychological well-being scores is not meaningful. However, when comparing Denmark and Sweden, very similar data structures, and thus comparable concepts of well-being, are found. Therefore, inter-country comparisons are warranted for these countries.

ARTICLE HISTORY

Received 11 September 2018
Accepted 15 May 2020

KEYWORDS

Principal component analysis; exploratory data analysis; covariance matrix; data structure

2010 MATHEMATICS SUBJECT CLASSIFICATIONS


62P15; 62H25

1. Introduction

Data comparability is a recurring topic in a wide array of statistical applications. Often, data will be collected in such a way that it is essentially divided into several subsets whose comparability needs to be assessed empirically. This happens for instance when data are collected across several centers (or countries) or when different versions of an instrument, the mode of administration of a survey, or the order of survey questionnaire items are applied. While a mixture of modes of administration, for example mail and telephone, can improve response rates in surveys, it can also be quite problematic by inducing differences in response behavior that may lead to biased results [3,14]. Powers, Mishra, and Young report effects of mode of administration on changes in mental health scores that are of a magnitude that is considered to be clinically meaningful [16], so the comparability issue

CONTACT Anne Helby Petersen  ahpe@sund.ku.dk  Department of Public Health, University of Copenhagen, Øster Farimagsgade 5, Copenhagen, Denmark

*Note: An early draft of this article was included in the conference proceedings *Symposium i Anvendt Statistik 2018. Copenhagen: The National Research Center for the Working Environment in Denmark & University of Copenhagen (2018)*.

 Supplemental data for this article can be accessed here. <https://doi.org/10.1080/02664763.2020.1773772>

is not just a statistical puzzle. Similarly, combining data obtained from different sampling schemes can also be problematic, as illustrated by Liu, who warns against combining online panel data with intercept samples (a pool of respondents obtained through banners, ads, or promotions) [12]. But mixing several data collection methods is essential for achieving further advances in surveying- and data collection methodology, as it allows us to simultaneously build on existing methods and answer new empirical questions. However, if we do not address the question of comparability among the data subsets arriving from e.g. different countries, we risk conducting analyses whose most fundamental assumptions are not satisfied.

The comparability issue in these examples can be summarized as follows: Assume that we have two datasets with the same variables, but different observations, often represented as a single dataset with a subset-inducing variable, and that we wish to compare them without specifying a model, or even a variable of interest. The central question is then whether the two datasets can readily be combined for the purpose of later data analysis, or if the subset-inducing variable implies heterogeneity that must be dealt with in later statistical modeling.

Sophisticated methods for addressing this question are available when we are willing to assume a statistical model. But this places the effort of assessing data comparability very late in the data analysis process and it makes the comparability assessment ad-hoc, as it essentially relies on the appropriateness of the modeling choices, which again depends on the structure of the data. The use of (parametric) models does thus not constitute a general data structure comparison method, but rather a fitted-model comparison method. It addresses the interplay between the model and the data, not the data alone. However, useful tools for initial, exploratory investigations of data comparability are mostly absent. What is needed are procedures that compare differences in data structures in two subsets of a dataset without assuming neither directional nor hierarchical relationships between the variables. Simple methods like variable-by-variable tests of distributional differences suffer from the drawback that they only address marginal differences and ignore the interplay between variables, but entry-by-entry comparisons of two empirical correlation matrices quickly become unmanageable as the number of variables increase.

We propose three visual tools for data structure comparisons, which we will refer to collectively as Principal Component Analysis-based Data Structure Comparisons (PCADSC). These methods use principal component decomposition of the empirical covariance matrix in the two subsets to create intuitive visualizations of data structure differences. The use of PCA implies that we can only compare datasets with numerical variables and that we focus on the linear aspect of the data structures. The proposed tools are implemented in the R package PCADSC [15].

1.1. Comparing psychological well-being

An example of where the data structure comparability issue must be considered is when conducting international rankings of countries according to e.g. educational quality (e.g. the PISA project) and citizen happiness (e.g. the UN *World Happiness Report* project). From a methodological point of view, such international rankings are problematic, as they rely on the fundamental assumption that the measured concepts are inherently the same across countries.

We use data from the 2012 version of the European Social Survey (ESS) project to investigate inter-country differences in psychological well-being and happiness. In the rankings of happiness, not much work has yet been devoted to evaluating the assumption of international comparability, though Veenhoven presents a theoretically thorough, but empirically simplistic, summary of possible reasons for lack of comparability [18], while Lolle & Andersen show highly potent translation issues for the term *happiness* [13]. If for instance two countries differ in terms of how social networks are typically built and structured, with one emphasizing family relations and the other mostly focusing on other social relations, then having a weak family connection does not have the same implications in the first country as it does in the second one. More specifically, whereas in the first country, lack of familial network might be related to loneliness, lack of general social capital, and isolation, in the second country, the quality of the family network might not be informative at all about other aspects of a person's social or psychological well-being. The two countries thus differ in how different aspects or measures of psychological well-being are interrelated, which is essentially a difference in data structures. And therefore, comparing the two countries in these measures is not a meaningful endeavor.

Our starting point is Denmark, a small, Northern European country that has repeatedly been awarded with the title of 'happiest country in the world' by the *World Happiness Report*, most recently in 2016 [6], and we wish to investigate if this title is really meaningful at all. In order to do this, we compare the Danish ESS psychological well-being data with that of Bulgaria. Though both countries are European and thus neither very far apart geographically nor culturally, these two countries have previously been highlighted to be very different in terms of what defines happiness [7]. Moreover, intra-European, regional differences in the relationship between social capital and happiness have also been demonstrated [17]. When comparing Northern European countries to other European countries, a much less pronounced relationship between the two concepts is found. In particular, interpersonal relations play a less important role in Denmark, compared to Bulgaria. Therefore, a successful method for data comparisons should be able to detect these differences by looking at data on psychological well-being from these two countries. We also compare the Danish data with Swedish data in order to investigate if the PCADSC tools actually do hold discriminatory power. Denmark and Sweden are both Northern European countries and are often deemed very similar in terms of culture and history. Therefore, we expect fundamental concepts, such as psychological well-being, to be similar across these two countries.

2. Methodology

PCADSC compares the covariance matrices of two subsets of a dataset. If all variables in the subsets are jointly normal with known means, the covariance matrices are sufficient statistics describing the joint distribution of all the variables, but even without the normality assumption, pairwise correlations and marginal variable variances are still interesting quantities describing linear interrelations between variables. This makes the empirical covariance matrix a reasonable place to start looking for differences in data structures. Computing empirical covariance matrices for the two subsets and comparing them entry-by-entry becomes increasingly difficult as the number of variables increases and for this

reason we decompose and recombine the covariance matrices using *principal component analysis* (PCA) to get an overview of differences between the two subsets.

We assume that the variables have been standardized within each data subset before PCA is conducted so that all variables have mean zero and unit standard deviation. This is equivalent to working with the empirical correlation matrices rather than the empirical covariance matrices. We also assume that all considered variables have a numerical interpretation, for example by being continuous or ordinal categorical. Moreover, we use the following notation: $\mathbf{X} \in \mathbb{R}^{n_x \times d}$ and $\mathbf{Y} \in \mathbb{R}^{n_y \times d}$ are datasets containing the same number of variables, d , but possibly different numbers of observations, n_x and n_y . We use X_j to refer to the j th variables of \mathbf{X} and x_{ij} to refer to the i th observation within that variable, while x_i is the full i th observation row. Note that $X_j \in \mathbb{R}^{n_x}$, $x_{ij} \in \mathbb{R}$, and $x_i \in \mathbb{R}^d$. We let $\bar{X} = (\bar{X}_1, \dots, \bar{X}_d)^T = (1/n_x) \sum_{i=1}^{n_x} x_i$ denote the variable averages and use $S_x = 1/(n_x - 1) \sum_{i=1}^{n_x} (x_i - \bar{X})(x_i - \bar{X})^T \in \mathbb{R}^{d \times d}$ to denote the empirical covariance matrix of \mathbf{X} .

2.1. Using PCA for data structure comparisons

The tools presented here are all based on comparing the PCA results across two different datasets that contain the same variables, \mathbf{X} and \mathbf{Y} . Let

$$\mathbf{Z} = \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix} \in \mathbb{R}^{(n_x+n_y) \times d}$$

be the combined dataset. For each of these three datasets, we complete the following steps (here described for \mathbf{X} only):

- (1) Standardize each of the variables to have mean zero and unit standard deviation. Let $\tilde{\mathbf{X}}_j \in \mathbb{R}^{n_x \times d}$ be the standardized dataset.
- (2) Form the principal component analyses

$$S_x = \frac{1}{n_x - 1} \sum_{i=1}^{n_x} \tilde{x}_i \tilde{x}_i^T = \sum_{j=1}^d \lambda_j^x \eta_j^x (\eta_j^x)^T$$

thereby obtaining loadings η_j^x and eigenvalues λ_j^x for $j = 1, \dots, d$.

The hereby obtained PCA decompositions of the correlation matrices can then be compared. The standardization implies that the diagonal elements of S_x , S_y , and S_z all equal 1, and thus also that

$$\sum_{j=1}^d \lambda_j^x = \sum_{j=1}^d \lambda_j^y = \sum_{j=1}^d \lambda_j^z = d.$$

This identity will simplify some expressions below. We note that the sequence of loadings η_j and their associated eigenvalues λ_j yield a simultaneous description of the structure of the dataset for all approximating dimensions q . This implies that the loadings and the eigenvalues can be used to investigate the structure of the dataset without the need to decide on an approximating dimension, q , a priori.




Figure 1. Graphical representations of models used for simulating example data. Square nodes represent observed variables, while rounded nodes represent latent variables. Arrows are used to illustrate causal links.

We present three diagnostic plots that are designed to shine a light on different types and levels of data structure differences, namely the *cumulative eigenvalue (CE) plot*, the *angle plot*, and the *chroma plot*. These plots are described in turn below. For the deepest understanding of the data structure differences in two datasets, we suggest using all of the three plots in the same order as they are presented. While we describe the three plot types from a purely theoretical point of view below, we also provide example illustrations as a supplement to the somewhat technical definitions. These plots are available in Figures 2–4 for the CE-, angle-, and chroma plots, respectively, and they are based on two simulated datasets:

- Dataset A:** This dataset contains 1000 independent simulations from the same underlying model, namely *model 1* from Figure 1. The observations are randomly divided into two groups.
- Dataset B:** This dataset contains 500 independent simulations from *model 1* from Figure 1 and 500 independent simulations from *model 2* in the same figure. The observations are naturally divided into groups corresponding to which model they were simulated from.

In *model 1* the observed variables V_3 and V_4 are caused by a common, unobserved variable L_1 . Further details on the models and the simulated data are available in the supplementary material, including covariance matrices for the two models, in Appendix A. In dataset A, all observations stem from the same, underlying data structure and thus, the PCADSC plots should point towards no notable differences in data structures. In dataset B, on the other hand, there are two fundamentally different, underlying data structures and therefore, the PCADSC plots should illustrate this lack of homogeneity in data structures.

2.2. The cumulative eigenvalue plot

The cumulative eigenvalue (CE) plot compares the eigenvalues of the correlation matrices. These eigenvalues represent how much information is withheld in each component

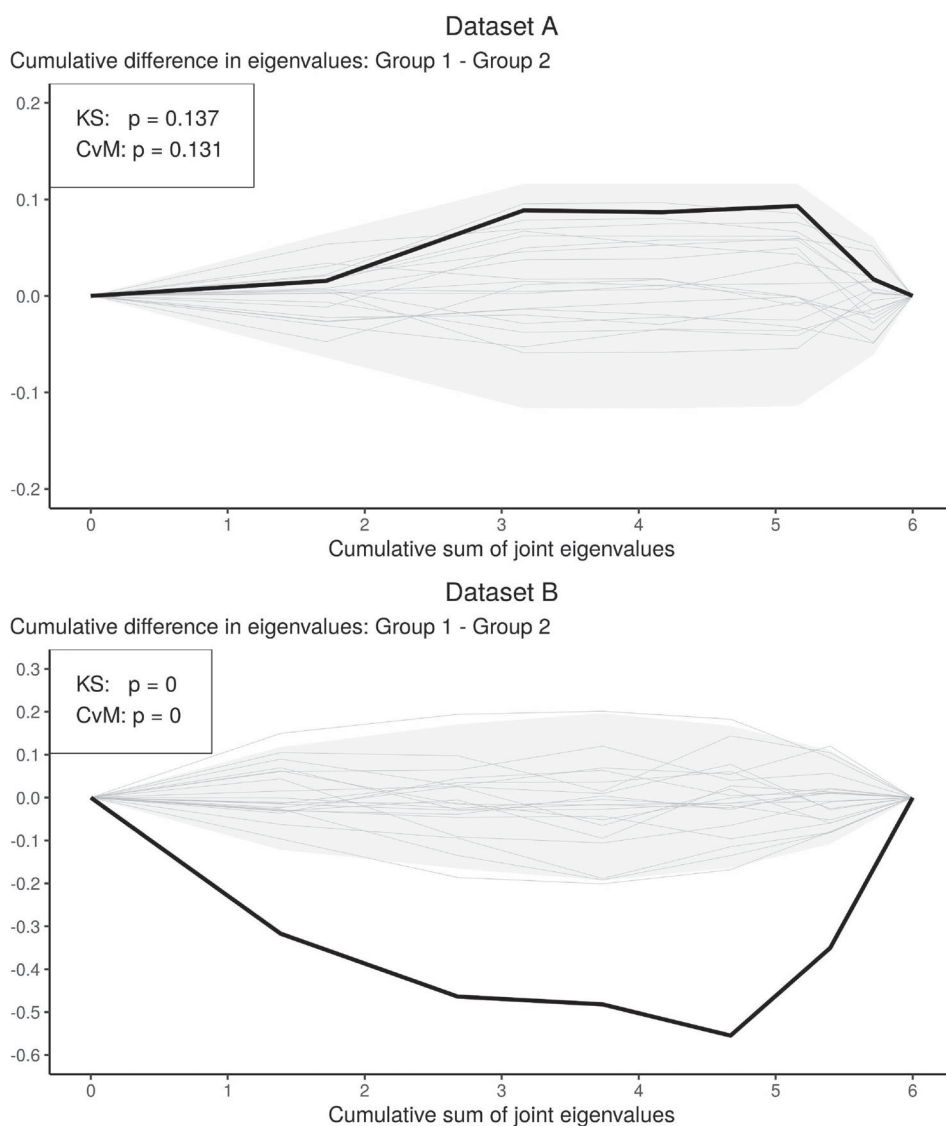


Figure 2. CE plots for the simulated datasets. Dataset A (with no differences in data structure) in top panel, and dataset B (where observations stem from different models) in bottom panel. Plots are annotated with the p -values of the Kolmogorov-Smirnov (KS) and the Cramér-von Mises (CvM) tests of the null hypothesis of no difference in data structures. The bold, black lines illustrate the observed cumulative eigenvalue differences, while the shaded areas are 95% confidence bands under the null hypothesis of no difference in data structures.

in terms of explained variance. Thus, by comparing cumulative sums of eigenvalues, it is possible to obtain a detailed picture of how the two datasets differ in terms of what components are the most informative. In order to investigate whether the same proportion of the total variation can be described by the same number of principal components in the

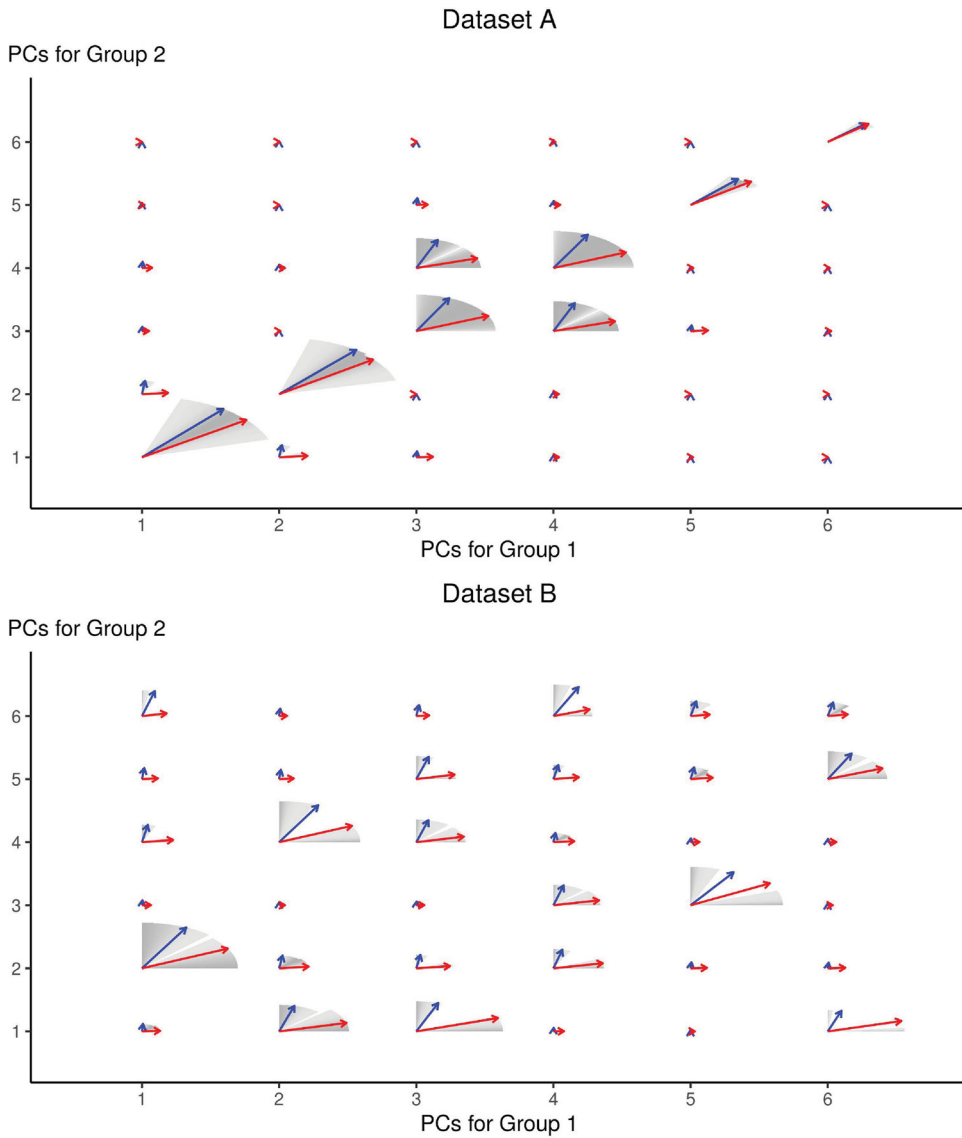


Figure 3. Angle plots for the simulated datasets. Dataset A in top panel and dataset B in bottom panel. Blue arrows show the principal components of the observations in group 1 decomposed in the coordinate system of the principal components of group 2, while the red arrows illustrate the reverse (Colour online).

two datasets, we plot a piecewise linear curve connecting the points

$$(0, 0), \quad (\lambda_1^z, \lambda_1^x - \lambda_1^y), \quad (\lambda_1^z + \lambda_2^z, \lambda_1^x + \lambda_2^x - \lambda_1^y - \lambda_2^y), \dots, \\ \left(\sum_{j=1}^d \lambda_j^z, \sum_{j=1}^d \lambda_j^x - \sum_{j=1}^d \lambda_j^y \right)$$




Figure 4. Chroma plots for the simulated datasets. Dataset **A** in top panel and dataset **B** in bottom panel. Each bar represents a principal component and the bars are annotated with their cumulative percentages of explained variance.

Due to the standardization, the last point will always be equal to $(d, 0)$. Thus, this curve will begin and end at the x-axis. And the larger excursions it makes away from the first axis, the less alike the cumulative eigenvalue sums for the two datasets are. Moreover, a positive cumulative difference implies that dataset **X** holds more information in the first components than dataset **Y** does. By using the cumulative eigenvalues of **Z** as the first coordinate, we obtain a visualization where a component take up an amount of horizontal space that corresponds to the variance explained by the component (that is, its eigenvalue), thereby letting the more influential components drive the visual impression.

In order to test whether these cumulative differences are statistical artefacts or if they represent something real, we have implemented both the *Kolmogorov-Smirnov* and the *Cramér-von Mises* test statistics, which are given by

$$KS = \max_{k=1,\dots,d} \left| \sum_{j=1}^k \lambda_j^x - \sum_{j=1}^k \lambda_j^y \right|, \quad CvM = \sum_{k=1}^{d-1} \frac{\lambda_k^z + \lambda_{k+1}^z}{2} \left(\sum_{j=1}^k \lambda_j^x - \sum_{j=1}^k \lambda_j^y \right)^2.$$

We conduct the tests as *permutation tests*, that is, by randomly reallocating the combined and separately standardized datasets into two new datasets of n_x and n_y observations, respectively, and then redoing the CE plot steps and recalculating the test statistics. This should be done a large (e.g. 10,000) number of times. Then, a p -value is obtained by computing the proportion of reallocated datasets that lead to test statistics at least as large as the one we found for the original datasets.

The permutation test results are also used to visualize the uncertainty of the CE curve in the plots. In the CE plots shown in the following section, we plot the observed curve together with 20 of the resampled curves, as well as a shaded region visualizing pointwise 95% coverage intervals. If the observed curve is very different from the resampled curves or if it is substantially outside the shaded region, then this also indicates differences between the two datasets.

Two examples of cumulative eigenvalue plots are available in Figure 2. In the top panel, we see a cumulative eigenvalue curve that falls well within the confidence region, thereby suggesting no difference in eigenvalues (as expected). This also agrees with the large p -values and the fact that only a single underlying model was used for generating this data. In the bottom panel, the conclusion is reversed, which is also in accordance with the nature of the simulated data, which contains two distinct groups.

2.3. The angle plot

The angle plot compares both loadings and eigenvalues at once and it can be used to understand the information loss if the data structure of one dataset is superimposed on the other, thereby revealing which principal components (i.e. loading and eigenvalue pairs) are most similar and most different across the two datasets. Let $\lambda_{\max} = \max\{\lambda_1^x, \lambda_1^y\}$ be the largest eigenvalue for the two datasets. Then the empirical correlation matrix for the \mathbf{X} dataset, S_x , has the following orthogonal decomposition in the coordinate system of the \mathbf{Y} dataset

$$S_x = \sum_{j=1}^d \lambda_j^x \eta_j^x (\eta_j^x)^\top = \lambda_{\max} \sum_{k=1}^d \left(\sum_{j=1}^d \sqrt{\frac{\lambda_k^x}{\lambda_{\max}}} \eta_j^y (\eta_j^y)^\top \eta_k^x \right) \left(\sum_{j=1}^d \sqrt{\frac{\lambda_k^x}{\lambda_{\max}}} \eta_j^y (\eta_j^y)^\top \eta_k^x \right)^\top,$$

and we have a similar decomposition of S_y in the coordinate system of the \mathbf{X} dataset. We propose to visualize these two decompositions in a $d \times d$ grid display. In the j th row and k th column of this display we plot two arrows based at the lower left corner of the grid cell. The first arrow has length μ_{jk} and angle $\theta_{jk}/2$ counterclockwise from the diagonal, and the second arrow has length ν_{jk} and angle $\theta_{jk}/2$ clockwise from the diagonal. To facilitate the following description we will refer to the arrows drawn counterclockwise as the blue

arrows, and the arrows drawn clockwise as the red arrows. The lengths μ_{jk} and ν_{jk} , and the angle θ_{jk} , are given by

$$\mu_{jk} = \sqrt{\frac{\lambda_k^x}{\lambda_{\max}}} |\eta_k^{x\top} \eta_j^y|, \quad \nu_{jk} = \sqrt{\frac{\lambda_j^y}{\lambda_{\max}}} |\eta_j^{y\top} \eta_k^x|, \quad \theta_{jk} = \arccos(|\eta_k^{x\top} \eta_j^y|).$$

For two d -dimensional, unit length vectors a and b , it holds that $a^\top b = \langle a, b \rangle = \tilde{c}(a, b)$, where \tilde{c} denotes the empirical correlation function. Thus, in the angle plot, we are essentially looking at the absolute values of correlations between loadings that have been scaled according to their variance contributions. The absolute value of the projection $\eta_k^{x\top} \eta_j^y$ is inserted due to the indeterminacy of the direction of loading vectors. This indeterminacy implies that the angle between loadings from the two datasets can always be chosen to be in the interval $[0, \pi/2]$, and hence the decomposition of S_x and S_y can be visualized in a joint plot by dividing the angles by two and using counterclockwise and clockwise shifts from the diagonal. Furthermore, the scaling of the lengths by λ_{\max} is made so that the longest arrow has at most unit length.

In the *angle plot*, the blue arrows in the k th column of the grid display visualize the decomposition of the k th principal component of the **A** dataset in the coordinate system of the **B** dataset. Similarly, the red arrows in the j th row visualize the decomposition of the j th principal components of the **B** dataset in the coordinate system of the **A** dataset. If we have coinciding blue and red arrows along the diagonal in the grid display, and arrows of vanishing lengths in the off-diagonal cells, then the structures of the two datasets are identical. In contrast, differences in the eigenvalues are visualized as differences in the lengths of the blue and the red arrows, also in the diagonal. And loadings in other directions than the corresponding loading from the other dataset are visualized as angle separation of the blue and the red arrows in the diagonal cells, as well as arrows of non vanishing length in the off-diagonal cells.

However, there is a generic situation where the angles between the arrows lack discriminative power, namely when two or more neighboring eigenvalues are identical or close. We refer to this situation as *eigenvalue indeterminacy*. For instance, multiple eigenvalues close to 1 arise if more than one of the variables are uncorrelated, or close to uncorrelated, of all the other variables. In this situation only the corresponding eigenspace is well estimated, but the rotation of the associated eigenvectors will be subject to sampling variability.

To investigate if we are in such a situation of identical data structures with neighboring eigenvalues being close, where comparably long arrows may occur in the corresponding blocks of off-diagonal cells, we have implemented a permutation test similar as done for the CE-plot. For each random reallocation we recompute the angles θ_{jk} . The resampled angles may both be used for quantitative tests, where large angles are critical for the diagonal cells and small angles are critical for the off-diagonal cells, as well as for the visualization of the sampled angle regions. In the visualization we propose to present the entire sampled region of angles for each cell, but gray shaded according to percentiles such that more critical angles are lighter gray. In our suggested visualizations, we do not present the p -values from each of the quantitative tests, but the p -values are available in the PCADSC R package [15].

Note that we can compute d^2 separate, but dependent, p -values. Therefore, inference drawn with reference to all of them will suffer from multiple testing issues and for that reason, we have devised a global test of difference in loadings. We use a Kolmogorov-Smirnov

type test statistic, namely

$$KS_{\text{angle}} = \max_{k=1, \dots, d^2} \log \frac{\text{median}_{(k)}}{p_{(k)}^{\text{obs}}},$$

where $p_{(k)}^{\text{obs}}$ is the k 'th smallest of the observed p -values and $\text{median}_{(k)}$ denotes the median over the k 'th smallest p -values, both for $k = 1, \dots, d^2$. Discrepancy from the null hypothesis of no difference in loadings will be manifested as $p_{(k)}^{\text{obs}} \ll \text{median}_{(k)}$ for some k , and thus the test statistic can be evaluated as a one-sided test with large values being critical. A global p -value for loading differences is then computed as a permutation test, similar to the procedure for the CE plot.

Figure 3 presents two examples of angle plots based on simulated data. For dataset **A**, we find very short off-diagonal arrows except for the off-diagonal cells corresponding to the 3rd and 4th eigenvalues. But since the arrows in cells (3, 4) and (4, 3) lie well within the shaded angle regions (supported by local p -values 0.9119 and 0.9103) this is in agreement with identical data structures, where the 3rd and 4th eigenvalues suffer from indeterminacy (in this situation due to the two independent variables in *model 1*). This is also supported by a global p -value from the Kolmogorov-Smirnov angle test of $p = 0.9988$. For dataset **B**, on the other hand, we find very little agreement in all components and a global p -value of 0.0000, which should be expected, as the two groups in the dataset correspond to different data structures.

2.4. The chroma plot

The chroma plot is primarily an illustration of differences in loading patterns and it targets the question of how the roles of the original variables are different between the two datasets, thus leading the data structure comparison question back to its original, empirical context. The chroma plot consists of two panels, one for each dataset, made up of colored bars. These bars each represent a principal component and their coloring illustrates the relative weights of the d original variables, that is, their absolute, normalized loading contributions. More specifically, when illustrating the i th principal component in the two data subsets, we plot vertical bars of length one that has been divided into d segments of different colors. The width of the j th colored segment is given by

$$\omega_{ij}^x = \frac{|\eta_{ij}^x|}{\sum_{k=1}^d |\eta_{ik}^x|}, \quad \omega_{ij}^y = \frac{|\eta_{ij}^y|}{\sum_{k=1}^d |\eta_{ik}^y|},$$

where η_{ij}^x and η_{ij}^y denotes the j th entry of η_i^x and η_i^y , respectively. Due to the indeterminacy of the loading signs, all the signs are removed from the coefficients in the loadings. The bars are ordered according to the eigenvalues and they are annotated with the cumulative percentage explained variance of that component, that is, the scaled and summed variance contributions,

$$\tilde{\Sigma}_i^x = \frac{\sum_{j=1}^i \lambda_j^x}{\sum_{k=1}^d \lambda_k^x} = \frac{\sum_{j=1}^i \lambda_j^x}{d}, \quad \tilde{\Sigma}_i^y = \frac{\sum_{j=1}^i \lambda_j^y}{\sum_{k=1}^d \lambda_k^y} = \frac{\sum_{j=1}^i \lambda_j^y}{d}.$$

Especially when d is large, we recommend plotting only a select set of interesting principal components (for instance identified by use of the angle plot). In this scenario, the annotations should rather be the non-cumulative variance contributions, $\tilde{\sigma}_i^x = \lambda_i^x/d$ and $\tilde{\sigma}_i^y = \lambda_i^y/d$. If there is eigenvalue indeterminacy for neighboring components, as discussed above, the loading patterns for these components will not have a meaningful interpretation. Therefore, one should not compare the chroma bars for component that were identified to suffer from eigenvalue indeterminacy by the angle plot.

In Figure 4, two examples of chroma plots are available. For dataset **A**, we see very similar visual patterns across the two panels for components 1, 2, 5 and 6, corresponding well with the fact that both groups are drawn from the same underlying model. Components 3 and 4 do not have similar loading patterns across the two subsets of the data, but this is expected since we identified these components to suffer from eigenvalue indeterminacy in the angle plot. For dataset **B**, we see different color compositions for the bars in the left- and the right panel, illustrating that the data structures are not identical for the two groups. The plots resulting from this procedure should be inspected focusing on two properties: Similarities in loading patterns, which will correspond to similar visual impressions, and similarities in variance contributions. For each component, the loadings describe how influential the different variables are on that component. Therefore, the chroma plot allows us to make qualitative statements about the original datasets, such as ‘*variable 1 is generally more influential in subset A than it is in subset B*’, thereby helping us to understand where and why potential data structure differences are found.

2.5. A data structure comparison workflow

Finally, we turn to a few comments about how the three plots presented thus far can be used together. We suggest that the plots are applied in the order in which they were presented, thereby moving from overall assessments of similarity to explanations about the nature of potential data structure differences. This workflow is summarized in Figure 5 and should follow three steps:

- (1) Make a CE plot. If this plots suggests similarity in eigenvalues then it is appropriate to move on to the angle plot to assess similarity in loadings. Otherwise, it can be concluded that the structures of the two datasets are not similar.
- (2) Make an angle plot. If this plots suggests similarities in both eigenvalues and loadings, then the data structure comparison exploration is concluded. If the angle plots suggests differences in data structures, identify which components are responsible for the differences.
- (3) Make a chroma plot of the components selected to be different by the angle plot, but whose differences are not due to eigenvalue indeterminacy among neighboring components. Use this information to discuss *how* and *why* the data structures might be different.

We will now turn to a data example, showing how the PCADSC tools can be used in practice.

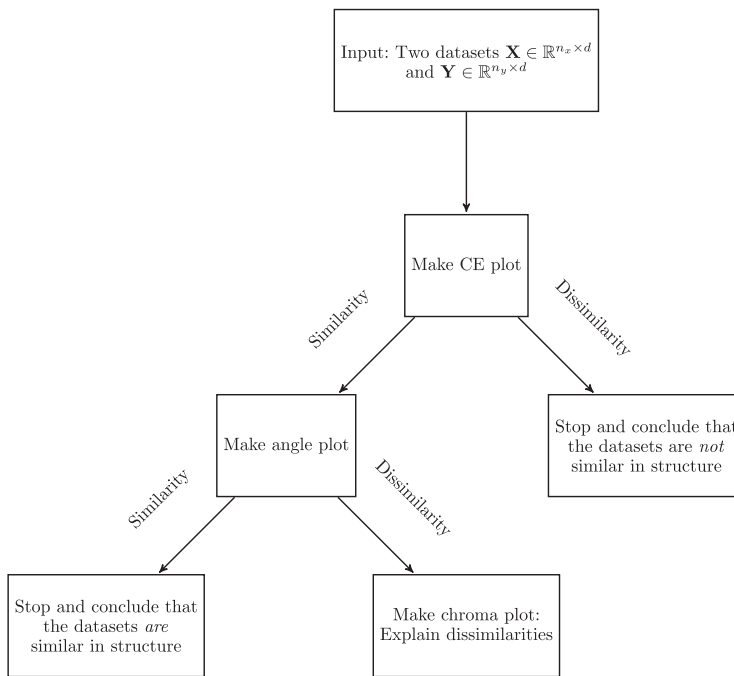


Figure 5. The PCADSC workflow, starting with two datasets and then successively applying the three visual tools for obtaining a deeper and deeper understanding of similarities and dissimilarities in the data structures.

3. Application

We now illustrate the three PCADSC tools using data from the 2012 version of the European Social Survey (ESS) project to investigate inter-country differences in psychological well-being and happiness. All computations and figures presented in this section were created using the R package PCADSC [15].

3.1. Data

The ESS 2012 data contain a total of 626 variables collected from 54,673 citizens of 29 countries. Here, we will only work with a subset of 35 questionnaire items that are all related to psychological well-being. These 35 items can be divided into 6 distinct scales, namely *Evaluative well-being*, *Emotional well-being*, *Functioning*, *Vitality*, *Community well-being*, and *Supportive relationships*. More details on these scales can be found in [7] and the relationship between questionnaire items and scales is summarized in Appendix B. We represent each of the scales by a single variable, which is calculated as the average score within the items related to that variable and scaled such that it takes a value between 0 and 10.

For simplicity, we use only complete cases for this construction and thus exclude all participants that did not answer all the 35 questionnaire items used below. An alternative approach is to use pairwise complete cases. We present results from PCADSC analyses using pairwise complete cases in Appendix C, but find no differences in conclusions relative to the results using complete cases (presented below). Note that both of these approaches

Table 1. Median values of each of the six dimensions of psychological well-being, stratified by country.

	Denmark	Bulgaria	Sweden
Evaluative well-being	8.75 (8.00, 9.50)	5.00 (3.50, 7.00)	8.00 (7.00, 9.00)
Emotional well-being	8.33 (7.22, 8.89)	6.67 (5.00, 7.78)	7.78 (6.67, 8.89)
Functioning	7.57 (6.93, 8.21)	6.68 (5.50, 7.68)	7.04 (6.39, 7.68)
Vitality	7.50 (6.67, 8.33)	7.50 (5.83, 8.33)	7.50 (6.67, 9.17)
Community well-being	6.77 (5.83, 7.57)	4.67 (3.70, 5.70)	6.57 (5.66, 7.37)
Supportive relationships	8.25 (7.42, 8.92)	7.25 (6.17, 8.08)	8.25 (7.42, 8.75)

Note: 25 and 75 percentiles are listed in parentheses. The scales are constructed such that they all run from 0 to 10.

for handling missing information are only consistent if the data are missing completely at random [11].

Approximately 9% of the observations in the Danish sample have one or more missing items, and the corresponding numbers are 20% in the Bulgarian sample and only 6% in the Swedish sample. All in all, we have $n_{DK} = 1498$ complete case observations in the Danish sample, $n_{BG} = 1798$ observations in the Bulgarian sample and $n_{SE} = 1736$ Swedish observations. Figure 1 summarizes the marginal distributions of the six dimensions of psychological well-being, stratified by country.


3.2. Comparing Denmark and Bulgaria

Figure 6 presents the CE plot and the angle plot obtained from comparing the Danish and Bulgarian psychological well-being scales. The CE plot show a remarkable degree of lacking comparability: The cumulative differences in the eigenvalues by far exceed what could come about randomly if there really were no difference in the data structures. This is also confirmed by the Kolmogorov-Smirnov and the Cramér-von Mises tests, which both result in p -values that are virtually zero. The CE curve generally lies above zero, suggesting that the Bulgarian dataset has larger eigenvalues for all but the last component, thereby explaining more variance at lower dimensionalities.

Moving on to the angle plot, we find that the differences are primarily to be found in the second, third, and fourth principal components (PCs). The blue arrows visualize the decomposition of the principal components for the Bulgarian dataset in the coordinate system of the Danish dataset. We see that PC2 also loads on PC3, that PC3 also loads on PC4, and that PC4 also loads on PC2 and PC3. The red arrows visualize the decomposition of the principal components for the Danish dataset in the coordinate system of the Bulgarian dataset. Here, we see that PC2 also loads on PC4, that PC3 also loads on PC2 and PC4, and that PC4 also loads on PC3. Thus, if we wish to understand why differences in the data structures occur, an inspection of the loadings of components 2, 3 and 4 might be informative. The global test of differences in loadings results in $p = 0.0000$, which supports the conclusion of differences in loading structures.

The chroma plot in Figure 7 allows us to look closer into these components. Here, we find that the relative importance of the *Community well-being* and *Supportive relationships* scales is much larger in the Bulgarian sample than in the Danish. In the Danish data, on the other hand, we find that *Vitality* and *Emotional well-being* seem to play bigger roles, as

Cumulative difference in eigenvalues: BG - DK



PCs for DK




Figure 6. CE plots (top panel) and angle plots (bottom panel) comparing Danish and Bulgarian data on psychological well-being. The CE plot is annotated with the p -values of the Kolmogorov-Smirnov and the Cramér-von Mises tests of the hypothesis of no difference in data structures. In the angle plot, the blue arrows show the principal components of the Bulgarian dataset decomposed in the coordinate system of the principal components of the Danish dataset, while the red arrows illustrate the reverse (Colour online).

they appear with larger loadings in more high-ranking components in this sample, relative to the Bulgarian.

All in all, we find that psychological well-being does *not* seem to be the same concept in Bulgaria and Denmark. The two countries disagree both in how many dimensions are needed to capture the most important parts of the concept (as illustrated by the differences in eigenvalues) and in how these dimensions are then weighted among the 6 scales (as

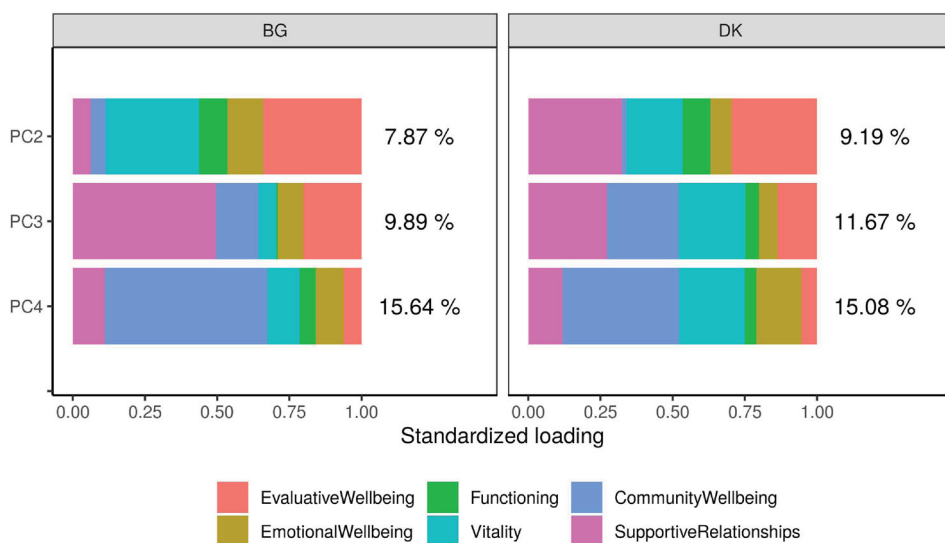


Figure 7. Chroma plot for comparing Danish and Bulgarian data on psychological well-being. A chroma plot comparing the 2nd, 3rd, and 4th principal components of the Bulgarian- and Danish psychological well-being data. The component-bars are annotated with their relative variance contributions.

illustrated by the angle- and chroma plots). In Bulgaria, interpersonal features seem to be more informative of psychological well-being, whereas in Denmark, individual characteristics play a relatively larger role, corresponding well with the previous findings mentioned above. Thus, the datasets are fundamentally different and we should therefore be wary about combining them in a joint analysis, which was also the conclusion of the ESS authors, though based on country-level aggregated statistics [7]. Moreover, this also implies that two countries cannot be ranked in terms of which country is ‘the most happy’, at least not by referring to psychological well-being dimensions such as those encountered here.


3.3. Comparing Denmark and Sweden

We now turn to the comparison of Denmark and Sweden in terms of psychological well-being. Figure 8 shows the CE- and angle plots for these two countries. In the CE plot, we now find the cumulative eigenvalue curve to be just within the acceptance region of the null-hypothesis. This is also reflected by the two tests, which now produce p -values of $p_{KS} = 0.17$ and $p_{CVM} = 0.11$, respectively. Thus, it is not unreasonable to assume equal eigenvalues in the two datasets.

The angle plot in Figure 8 shows that the two datasets agree very strongly about the relative importance of the six scales in the six PCs, as almost all off-diagonal arrows are practically non-existent. This is furthermore supported by the global test, which results in $p = 0.1980$, also suggesting no difference in loadings. This implies that if one already has the information held in the first PC from the Danish data, this information is in itself mostly sufficient to describe the first PC of the Swedish data.

Looking at the chroma plot in Figure 9, the same tale is told once again: Here, we find remarkably similar loading patterns in the first three components (which are responsible

Cumulative difference in eigenvalues: DK - SE



PCs for SE




Figure 8. A CE plot (top panel) and an angle plot (bottom panel) comparing Danish and Swedish data on psychological well-being. The blue arrows show the principal components of the Danish dataset decomposed in the coordinate system of the principal components of the Swedish dataset, and the red arrows illustrate the reverse (Colour online).

for almost 80% of the variance in both datasets), and slight, but increasing, differences in the remaining three components. We therefore conclude that any differences in the data structures of the Danish and the Swedish samples are related to the least important dimensions of the datasets and that these dimensions are only responsible for less than 25% of the variance in both datasets. In particular, this means that we can combine and compare the Danish and Swedish datasets in a meaningful way and for instance conclude using Figure 1 that in general, Danes seem to be somewhat more happy than Swedes, and in particular that the least happy people in Denmark (represented by the 1st quartiles) are generally happier




Figure 9. Chroma plot comparing the loading patterns of the Danish and the Swedish subsamples. For each component, the bar is annotated with its cumulative variance contribution, that is, how much variance can be explained by having information of this and the preceding components.

than the least happy people in Sweden. A more thorough, statistical investigation could now be put to work on answering *why* this seems to be the case.

4. Conclusion and discussion

Three new tools, referred to collectively as Principal Component Analysis-based Data Structure Comparisons (PCADSC), for the task of deciding if two datasets can be combined for analysis were proposed and discussed. They all employ the principal component decomposition of the empirical covariance matrix performed on two subsets of a dataset in order to create three intuitive visualizations of data structure differences. The first visualization, the cumulative eigenvalue plot, compares the eigenvalues from the two datasets. The second visualization, the angle plot, uses the interpretation of eigenvalues as explained variance to focus a comparison of loadings (eigenvectors). Finally, the third visualization, the chroma plot, allows for qualitative explanations of differences in loading patterns in selected principal components.

In an analysis of data from the 2012 version of the European Social Survey (ESS) project, we illustrated that the PCADSC tools help inform analyses of inter-country differences in psychological well-being and happiness. It should be noted that even though concerns about pooling data from different countries are not yet very prevalent when it comes to happiness rankings, the subject has been very well scrutinized in the field of international educational rankings. For instance, the PISA tests have repeatedly been criticized for not being meaningful objects of international comparisons, especially due to problems with differential item functioning [8,9,19] and translation problems [2]. This highlights lack of data comparability as an empirically relevant and concerning issue that there is no way around: When conducting rankings, the issue of data structure heterogeneity should always be addressed.

Usually, when using PCA, the main interest lies in the *scores*, i.e. the coordinates $\eta_j^\top (x_i - \bar{x})$ of the projections observations onto the loadings, but the eigenvalues and the accompanying loadings are in fact more suitable objects for assessing data structure similarity. If two dataset subsets have similar loadings, then they can be argued to measure the same underlying quantities. Similar loadings can occur in situations where the two sets of scores are different, for example if the two datasets come from two different populations of subjects. On the other hand, if the loading patterns are different in the two subsets, then this indicates that their respective variable interplays differ, and hence it would be criticizable to use these variables for comparisons across the two subsets. The PCADSC tools are novel in their emphasis on the information held in the loadings rather than the scores.

Though the PCADSC toolbox comprises three diagnostic plots as of now, the PCA decomposition can of course be utilized further to bring about new PCADSC plots that illustrate more aspects of data structure differences. For instance, the scree plot (as introduced by Cattell [4]), which plots the ordered eigenvalues against their component numbers, could easily be recreated as a PCADSC tool by creating scree plots for both data subsets in the same coordinate system, yielding a simple graphical comparison. The resampling approach used in the CE plot could then be adopted, thus annotating the scree plots with resampled curves. This general approach to model-fit evaluation is very helpful in combination with permutation tests, as inspired by Lin and colleagues [10].

The proposed methods also have limitations. First of all, as PCA is performed after standardization of the variables, the PCADSC methods cannot detect differences in neither mean values nor in marginal variances – this information is thrown away as the very first step. However, such differences are of a fundamentally different nature than those we have discussed so far. Differences in mean values are often the main interest of the analysis and should of course not be regarded as a data comparability issue. Differences in marginal variances, on the other hand, can pose some modeling challenges, but none that are not generally solvable by use of random effects models. We therefore do not consider this property of PCADSC to be a major drawback.

A more prominent limitation is the fact that the PCASDC methods are only valid for variables that have a numerical interpretation. In particular, this means that the methods cannot be used for data with nominal categorical variables. This limitation is inherited from the PCA method and thus, moving beyond it will entail replacing the PCA framework with a more general one. Multiple correspondence analysis (MCA) is often suggested as a categorical extension of PCA [1] and including MCA in the PCADSC tools is thus a natural next step.

Moreover, though we have illustrated the discriminatory power of PCASDC in one simulated- and on real-life data example, of course further evaluation of its performance is needed. While the best feature of PCADSC is perhaps the intuitive, visual nature of the tools, it is also the biggest weakness. For statistical methods that result in a number (or a set of numbers), the standard way of evaluating performance is to report on simulation studies, but in the case of statistical methods that report graphics this is not feasible. Therefore, a systematic evaluation of the performance of the PCADSC methods is not very straightforward to design: It really depends on whether or not users can detect actual differences in data structures by using the diagnostic plots. We suggest that the variability of the methods could be studied further using *Wally plots* [5], which are learning tools developed for teaching students how to interpret residual plots. First in line as a topic for more thorough

investigations of performance is the sensitivity towards the sample sizes of the two datasets, n_1 and n_2 .

A different aspect of the data size, namely the number of variables d , may also limit the useability of the PCADSC tools. While the CE plot scales well with large d , the angle plot and the chroma plot can become difficult to read if there are a lot of variables and thus a large number of principal components. However, because the angle plot lets the arrow length depend on the size of the corresponding eigenvalue, its visual impression should help guide a user towards influential *and* dissimilar components even though d is large. And when these components have been identified, the chroma plot can focus only on them, thereby avoiding plotting the loading patterns of too many components.

All in all, the PCADSC methods represent a first step towards addressing a known issue that has otherwise been met with unsatisfactory ad-hoc methods whose assumptions rest upon the very hypothesis they are testing. The PCADSC tools perform quite well in the data example provided here, yielding conclusions that are in accordance with well-established theories and they are simple to use and test further using the PCADSC R package.

Disclosure statement

No potential conflict of interest was reported by the author(s).

ORCID

Anne Helby Petersen  <http://orcid.org/0000-0003-3801-7230>

References

- [1] H. Abdi and L.J. Williams, *Principal component analysis*, Wiley Interdiscip. Rev. Comput. Stat. 2 (2010), pp. 433–459. Available at <http://dx.doi.org/10.1002/wics.101>.
- [2] M. Asil and G.T.L. Brown, *Comparing OECD PISA reading in English to other languages: Identifying potential sources of noninvariance*, Int. J. Test. 16 (2016), pp. 71–93.
- [3] D.J. Brambilla and S.M. McKinlay, *A comparison of responses to mailed questionnaires and telephone interviews in a mixed mode health survey*, Am. J. Epidemiol. 126 (1987), pp. 962–971. Available at <http://www.ncbi.nlm.nih.gov/pubmed/3661543>.
- [4] R.B. Cattell, *The scree test for the number of factors*, Multivariate Behav. Res. 1 (1966), pp. 245–276.
- [5] C.T. Ekstrøm, *Teaching 'instant experience' with graphical model validation techniques*, Teach. Stat. 36 (2014), pp. 23–26. Available at <http://doi.wiley.com/10.1111/test.12027>.
- [6] J. Helliwell, R. Layard, and J. Sachs, *World Happiness Report 2016 Update*, Sustainable Development Solutions Network, New York, 2016.
- [7] K. Jeffrey, S. Abdallah, and A. Quick, *Europeans' Personal and Social Wellbeing: Topline Results from Round 6 of the European Social Survey*, ESS Topline Results (Series 5), Centre for Comparative Social Surveys, London, 2015.
- [8] M. Kankaraš and G. Moors, *Analysis of cross-cultural comparability of PISA 2009 scores*, J. Cross. Cult. Psychol. 45 (2014), pp. 381–399.
- [9] S. Kreiner and K.B. Christensen, *Analyses of model fit and robustness. A new look at the PISA scaling model underlying ranking of countries according to reading literacy*, Psychometrika 79 (2014), pp. 210–231.
- [10] D.Y. Lin, L.J. Wei, and Z. Ying, *Model-checking techniques based on cumulative residuals*, Biometrics 58 (2002), pp. 1–12.
- [11] R.J.A. Little and D.B. Rubin, *Statistical Analysis with Missing Data*, 2nd ed., J. Wiley, New York, 2002, Chap. 3.

- [12] M. Liu, *Comparing data quality between online panel and intercept samples*, Methodological Innovations 9 (2016), p. 2059799116672877. Available at <http://dx.doi.org/10.1177/2059799116672877>.
- [13] H.L. Lolle and J.G. Andersen, *Measuring happiness and overall life satisfaction: A Danish survey experiment on the impact of language and translation problems*, J. Happiness Stud. 17 (2016), pp. 1337–1350. Available at <http://dx.doi.org/10.1007/s10902-015-9646-4>.
- [14] C.A. McHorney, M. Kosinski, and J.E. Ware, *Comparisons of the costs and quality of norms for the SF-36 health survey collected by mail versus telephone interview: Results from a national survey*, Med. Care. 32 (1994), pp. 551–67. Available at <http://www.ncbi.nlm.nih.gov/pubmed/8189774>.
- [15] A.H. Petersen and B. Markussen, *PCADSC: Tools for Principal Component Analysis-Based Data Structure Comparisons*, R package version 0.8.0, 2017. Available at <https://CRAN.R-project.org/package=PCADSC>.
- [16] J.R. Powers, G. Mishra, and A.F. Young, *Differences in mail and telephone responses to self-rated health: Use of multiple imputation in correcting for response bias*, Aust. N. Z. J. Public. Health. 29 (2005), pp. 149–154. Available at <http://www.ncbi.nlm.nih.gov/pubmed/15915619>.
- [17] A. Rodríguez-Pose and V. von Berlepsch, *Social capital and individual happiness in Europe*, J. Happiness Stud. 15 (2014), pp. 357–386.
- [18] R. Veenhoven, *Cross-national differences in happiness: Cultural measurement bias or effect of culture?*, Int. J. Wellbeing 2 (2012), pp. 333–353.
- [19] R.J. Zwitser, S.S.F. Glaser, and G. Maris, *Monitoring countries in a changing world: A new look at DIF in international surveys*, Psychometrika 82 (2017), pp. 210–232.